REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden to Washington Headquarters Service, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188) Washington, DC 20503.
PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.

1. REPORT DATE *(DD-MM-YYYY)*
OCTOBER 2012

2. REPORT TYPE
CONFERENCE PAPER (Post Print)

3. DATES COVERED *(From - To)*
NOV 2009 – JUN 2012

4. TITLE AND SUBTITLE
AN ARCHITECTURE TO SUPPORT INFORMATION AVAILABILITY IN THE TACTICAL DOMAIN

5a. CONTRACT NUMBER
IN-HOUSE

5b. GRANT NUMBER
N/A

5c. PROGRAM ELEMENT NUMBER
N/A

6. AUTHOR(S)
James Metzler, Brian Holmes, and Matthew Renodin

5d. PROJECT NUMBER
558S

5e. TASK NUMBER
IH

5f. WORK UNIT NUMBER
NC

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)
Air Force Research Laboratory/Information Directorate
Rome Research Site/RISA
525 Brooks Road
Rome NY 13441-4505

8. PERFORMING ORGANIZATION REPORT NUMBER
N/A

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)
Air Force Research Laboratory/Information Directorate
Rome Research Site/RISA
525 Brooks Road
Rome NY 13441-4505

10. SPONSOR/MONITOR'S ACRONYM(S)
N/A

11. SPONSORING/MONITORING AGENCY REPORT NUMBER
AFRL-RI-RS-TP-2012-062

12. DISTRIBUTION AVAILABILITY STATEMENT
APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED. PA Case Number: 88ABW-2012-4308
DATE CLEARED: 03 AUG 2012

13. SUPPLEMENTARY NOTES
Proceedings Military Communications (MILCOM) Conference, Orlando, FL. 30 Oct – 1 Nov 2012. The U. S. Government is joint owner of the work and has the right to copy, distribute, and use the work.

14. ABSTRACT
Access to information is critical to both the commander in an AOC and the warfighter in the field. Typically information is readily available at centralized command posts, however, at the tactical edge, resources are far more limited, making information dissemination a challenge. Targeting pods, already found mounted to the hardpoints of many tactical aircraft, provide a cost effective platform for making information available to tactical users. To this end, the Network-Centric Exploitation and Tracking (N-CET) program is designing, developing, and implementing a proof of concept architecture for pods that is net-centric, reconfigurable, and allows processing at the sensor. The approach taken to achieve these attributes is to embed processing and communications on the pod, and employ net-centric exploitation and fusion algorithms to distil information from high fidelity sensor data. Information Management services provide the interface between the sensors, processing, and network, disseminating information between algorithms, and prioritizing it as it goes out over the network. This paper provides an overview of the N-CET architecture and the sensors and net-centric algorithms integrated to evaluate the performance of the architecture through ground based experimentation.

15. SUBJECT TERMS
Information management, network-centric, distributed information systems, high performance computing

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
					JAMES METZLER
a. REPORT	b. ABSTRACT	c. THIS PAGE			19b. TELEPHONE NUMBER (*Include area code*)
U	U	U	UU	7	N/A

An Architecture to Support Information Availability in the Tactical Domain

James M. Metzler, Brian R. Holmes, and Matthew P. Renodin

Abstract—**Access to information is critical to both the commander in an AOC and the warfighter in the field. Typically information is readily available at centralized command posts, however, at the tactical edge, resources are far more limited, making information dissemination a challenge. Targeting pods, already found mounted to the hardpoints of many tactical aircraft, provide a cost effective platform for making information available to tactical users. To this end, the Network-Centric Exploitation and Tracking (N-CET) program is designing, developing, and implementing a proof of concept architecture for pods that is net-centric, reconfigurable, and allows processing at the sensor. The approach taken to achieve these attributes is to embed processing and communications on the pod, and employ net-centric exploitation and fusion algorithms to distil information from high fidelity sensor data. Information Management services provide the interface between the sensors, processing, and network, disseminating information between algorithms, and prioritizing it as it goes out over the network. This paper provides an overview of the N-CET architecture and the sensors and net-centric algorithms integrated to evaluate the performance of the architecture through ground based experimentation.**

Index Terms—**Distributed information systems, High performance computing, Information management**

I. INTRODUCTION

INFORMATION is of vital importance at all echelons of a military force. A commander requires information on the status of the forces under his or her control, the activity and intentions of the enemy, the location of civilians, and environmental conditions such as weather. This information can in most cases be made readily available to decision makers in centralized command posts with sufficient resources. Similar information is just as crucial to the warfighters carrying out the commander's intent, and in doing so, making decisions on their own. Unfortunately, at the tactical edge, resources are far more limited, making every literal bit of information critical, not only in that it is received, but also that there is value in it being sent.

Manuscript received August 6, 2012.

J. M. Metzler and B. R. Holmes are with the Air Force Research Laboratory, Information Directorate, Rome, NY, 13441 USA (phone: 315-330-4033; fax: 315-330-7009; e-mail: james.metzler@rl.af.mil, brian.holmes@rl.a.mil).

M. Renodin is with ITT Exelis, Rome, NY, 13441 USA (e-mail: matthew.renodin@exelisinc.com).

A. Information at the Tactical Edge

There are many reasons tactical warfighters need information. Most common is situational awareness (SA), the knowledge of the surrounding environment. In the USAF, a Joint Terminal Attack Controller (JTAC) is responsible for calling in close air support. Delivering ordinance safely on target requires SA of where the targets are, what assets are available, and the location of friendly forces.

Often a JTAC's responsibility extends to nominating targets, as in the case of discovering and identifying High Value Individuals (HVIs). This task requires a variety of intelligence products from sources such as ISR assets, intelligence agencies, other troops, and local civilians. Simply identifying a target is not sufficient for making the decision to prosecute it. Supporting evidence is necessary to verify the target and ensure collateral damage is minimized. Developing the target in this manner takes time and information – scarce resources in the fast pace of war at the tactical edge.

The availability of information at the tactical edge has several requirements. First, it must be delivered in near real-time despite potentially limited bandwidth resources. Additionally, communication between tactical users and assets should be made machine-to-machine to whatever extent possible to increase speed and reduce transcription error.

B. Tactical Information Dissemination Challenges

The tactical environment is characterized by unpredictability, scarce resources, and often a wealth of data with little or no context. Given these conditions, the right information is not always readily available to tactical users. For information to provide value, it must be relevant, accessible, and timely. When information is not relevant, not only is the bandwidth used to transmit that information wasted, but consumers can become overloaded with information. This can be distracting when a constant stream of full motion video (FMV) becomes a "face magnet", and counterproductive when important information is diluted by that which is not relevant.

The heterogeneity of today's military networks makes accessing information a significant challenge. Compounding this is the bandwidth limitations of military networks, especially at the tactical edge, and the fact that a tactical user may not know where information is, or that it even exists.

There are several causes for information not being delivered in a timely manner. Low bandwidth and high latency in military tactical networks can cause delays in the delivery of critical pieces of information. This problem is exacerbated by the introduction of less important packets to the network. Video, for example, can overwhelm a network and prevent time sensitive information such as target data from reaching its destination in time. Latency is also inherent in the processing of data, especially when that processing is conducted by an analyst at a remote location. In many cases, information would be timelier, and therefore more useful, going right from the sensor to the shooter.

While network technologies are advancing and bandwidth is increasing in both the tactical and strategic domains, the military's ability to collect data still far exceeds its ability to transmit it. Sensor resolution is growing at unprecedented rates, but the means to move the data being generated has not improved at a commensurate pace. The large datalinks required to accommodate high-bandwidth sensors, such as the CDL [1], require vast amounts of spectrum, and are not man-packable, so tactical users can be excluded from access to high fidelity sources. And when not enough bandwidth is available, data must be stored and processed post-mission, or in the worst case, dropped on the floor.

As the number and variety of sensors increase, so does the opportunity for them to collaborate on sensing by taking advantage of complimentary sensor modalities and varying geometries. However, the prevalence of stove-piped architectures often limits this ability, resulting in increased geolocation times and wasted sensor tasking.

Addressing these issues requires upgrading the sensors and avionics on current systems. Because this requires modification to an Operational Flight Program (OFP), these upgrades come at a high cost and have long timelines [2]. Upgrading an aircraft's capabilities through the addition of a sensor or targeting pod, such as Northrop Grumman's LITENING AT system [3], is appealing because doing so does not alter the OFP of an aircraft and therefore allows for quick integration and upgrade cycles.

C. Architecture Attributes

The research presented in this paper attempts to address the challenges of information availability at the tactical edge by designing, developing, and implementing a proof of concept architecture for pods that has the following attributes: 1) net-centric, 2) reconfigurable, and 3) allows processing at the sensor while preserving raw data. This is a continuation of the Network-Centric Exploitation and Tracking (N-CET) program initially presented in [4].

N-CET strives for net-centricity, that it may support multiple users and make information available in a timely manner. Tactical radios, organic to the pod, permit direct communication with tactical users. This shortens the path between sensor and shooter, both for the flow of sensor data to the user, and for requests for information from the user to the sensor. These radios also enhance communications between platforms, facilitating collaborative sensing and encouraging processing algorithms that utilize multiple sensing modalities and varying geometries.

The reconfigurability of the N-CET architecture allows it to support a variety of sensors and algorithms. This makes the architecture portable to different systems with varying hardware and software, but also allows for quick changes to a single system when the mission or environment dictates that a different sensor be used or alternate algorithm employed.

Key to providing timely, relevant information to tactical users is placing processing at the sensor as well as providing a means to archive raw sensor data. When data can be processed at the sensor, in its highest fidelity form, it can be exploited with minimum latency and bandwidth requirements, and contextualized for the user. This processing capability also provides a means for the information to be managed, so that it may be prioritized and disseminated to the users to which it is relevant, making the most of resource constrained tactical networks. Information that cannot be processed in real-time or sent off-board is archived so that it may be accessed later if necessary, such as at the occurrence of a significant event.

D. Outline

The remainder of this paper details the N-CET architecture and the results of experimentation. Section II presents the core elements of the architecture. Section III describes the sensors and algorithms that were integrated and used for experimentation, the results of which are presented in Section IV. Envisioned future work is discussed in Section V.

II. Core Architecture

The N-CET architecture is comprised of three core components: information management (IM), embedded processing, and communications. IM combined with organic communications provide net-centric capabilities that allow participants (sensor, algorithms, and users) to discover other participants, communicate, synchronize, and collaborate. IM also provides common interfaces between participants to facilitate reconfigurability. Embedded processing in the form of general purpose computing hardware supports data processing at the sensor and onboard storage provides archival capacity.

A cost effective method of providing these capabilities to tactical assets is through the use of pods. Whereas upgrading an aircraft to include new processing and communications is an expensive and lengthy process, placing these capabilities in a pod which the aircraft is already able to carry requires only power and a ride. This research has not limited the N-CET architecture to a pod form factor. In fact, the architecture lends itself well to many ISR systems. However, it is envisioned that a transition path will be found through the use of a pod such as the LITENING AT.

A. Information Management

For information to be relevant, available, and timely, a means for it to be contextualized, prioritized, and disseminated must exist. Without context, the value of one piece of information from another is indistinguishable. Context can come from the information itself, such as elements of time and location, and from processing it with exploitation and fusion algorithms. Based on context, *different* information can have distinct values to a *single* consumer, and the *same* information can have distinct values to *different* consumers. The value of the information can be used to prioritize it so that in the face of insufficient bandwidth to disseminate all information, the information of the highest value is preferentially treated. Of course, this can also ensure that information with no value is not disseminated at all.

The technologies developed by the Information Management research group at the AFRL Information Directorate provide the capabilities to contextualize, prioritize, and disseminate information through the use of publish, subscribe, and query. Information is encapsulated as a Managed Information Object (MIO) consisting of a payload (e.g. an image) and metadata (e.g. time and location), which provides context for the payload. Producers publish MIOs to the Information Management Services (IMS), which broker the MIOs based on the subscriptions registered by consumers and disseminate MIOs appropriately. MIOs are also archived by the IMS so that consumers may query the system for information that was produced in the past, at a time they may not have been connected or interested.

The publish/subscribe/query paradigm offers several benefits over typical point to point communication, perhaps the most significant being the decoupling of producers (e.g. sensors) and consumers (e.g. algorithms, users). It is not necessary for an information producer to be burdened with sending data to consumers that may be connecting and disconnecting intermittently. The producer simply publishes MIOs to the IMS so that they may be disseminated to consumers based on subscriptions. This decoupling facilitates reconfiguration because producers and consumers can be substituted. As long as they produce/consume the same MIO type, their replacement is transparent to other participants in the system.

The IM platform implemented in N-CET is the Phoenix Information Management Services [5]. Phoenix's SOA-based design separates IM tasks, such as submission, brokering, archival, and dissemination, into services that may be orchestrated, distributed, and substituted.

Phoenix is implemented as a server allowing clients to connect to service interfaces to publish, subscribe, and query information. Phoenix also facilitates the streaming of information directly from one client to another whereby the producer acts as a service providing information to the consuming client. This out-of-band delivery is beneficial in cases where a consumer wants all of a type of information generated by producer (e.g. FMV) because it does not incur the burden and latency of brokering each instance of information. Streaming utilizes the pub/sub methodology to manage information regarding the connection, allowing consumers to discover sources of the information they are interested in.

A cross-language interface is required to allow C++ clients to utilize the Phoenix services implemented in Java. Thrift generated source code provides an efficient means to create and manage this interface. Thrift is an open source software library and code-generator that provides a method for creating cross-language static code that communicates over a network [6]. In addition, Thrift was modified and extended to provide streaming capabilities between C++ and Java clients.

B. Embedded Processing

While Phoenix does not require an excessive amount of processing power, effectively managing information at high data rates requires processing capabilities not available in standard aircraft avionics. Beyond Phoenix, computing requirements are dictated by the class of sensors and algorithms on-board the pod. Whatever hardware is chosen, it must fit within the size, weight and power (SWAP) envelope of the form factor in which it is embedded. This research has used standard commercial hardware to develop a proof of concept, keeping in mind the SWAP requirements of a pod.

The processing hardware used by N-CET remains mostly unchanged from that which is described in [4]. Future development will likely focus solely on GP-GPUs as they are more power efficient than the Cell BE processor [7] and are available in a form factor suitable for the backplane of a pod. The core architecture also includes HDD storage for data that cannot be processed in real-time. Airborne applications will use solid state drives with read/write speeds exceeding that of HDDs.

C. Communications

To be net-centric and reconfigurable implies supporting a variety of datalinks which N-CET achieves by attempting to remain radio agnostic. Making the radio organic to the pod architecture also facilitates faster upgrades, allowing the pod to stay compatible with the warfighter's tactical radio, rather than the warfighter carrying an extra radio to match that of the platform. For testing purposes, COTS IP radios with representative data rates are used.

The transition to the Phoenix IMS has allowed N-CET to leverage two related technologies: QoS Enabled Dissemination (QED) [8] and Virtual Interface Approach to Cross-Layer Communications (VIA) [9]. QED is a set of QoS management services that, when integrated into the Phoenix IMS, monitor the QoS characteristics of the system (e.g. bandwidth and CPU) and manage the resources based on mission requirements defined by policy. Prior to QED, N-CET was only able to prioritize packets going into the network to influence the percentage of bandwidth that was allocated to a type of information. QED extends that capability to include managing which information is sent into the network. In addition, QED manages conflicting demands for resources as well as resource bottlenecks, and dynamically adjusts to changing mission requirements and resource availability.

For QED to be effective at managing bandwidth resources, it must have visibility into the performance of the underlying network. VIA provides this visibility by notifying QED of nodes affected by bottlenecks, allowing QED to throttle back transmission rather than saturating a node. VIA's Weighted Fair Queuing (WFQ) provides prioritized packets with a fair share of the available resources rather than a predetermined dedicated amount, ensuring no traffic flows are starved. VIA also provides capacity estimation to determine the available bandwidth that is allotted through WFQ.

III. Evaluating The Architecture

Sensors and algorithms were chosen that test the net-centricity of the architecture or benefit from the net-centricity that the architecture provides. The integration time of new sensors and algorithms was considered as well the ability of the architecture to support multiple (interchangeable) sensors and algorithms of the same class (e.g. FMV, GMTI tracking, video chipping). Sensors were also chosen based on their ability to generate high bandwidth data that stress the computing and storage capabilities of the architecture. Algorithms were chosen that could process this data into a product suitable for low bandwidth tactical datalinks. The sensors and algorithms chosen support the mission to detect, locate, identify, and track a target of interest.

A. *Sensors*

For the purposes of the mission above, in addition to the core components described in Section II, each N-CET node (pod surrogate) has an RF direction finding (DF) sensor and a high definition EO/FMV camera. Also, a ground moving target indicator (GMTI) radar has been simulated, supplying a sensor input from an external source.

B. *Algorithms*

Algorithms are integrated into N-CET as clients of the Phoenix IMS. Their interaction with data is through publication, subscription and query of MIO types registered in Phoenix. The algorithms integrated since [4] will be described below, and readers are directed to [4] for further details on existing algorithms.

The Controller is the nerve center of each N-CET node. It subscribes to *rfIntercepts* (containing the line of bearings (LOBs) to RF emitters) being generated by multiple nodes, triangulates the position of the target, and, depending on the mission, cross-cues the EO/FMV sensor onto the target. The N-CET architecture allows the Controller to receive this information from multiple nodes with the low latency necessary to image an RF emitter shortly after the radio is keyed. The architecture also allows the Controller to subscribe to other remote sources of targets, such as a GMTI platform publishing detections, and users publishing requests for imagery and video. The Controller tasks the EO/FMV sensor based on the mode of the system, for example, cross-cueing on RF and GMTI targets when in automated mode (while interleaving user imagery requests), and ignoring them when capturing FMV of a user defined target.

One appealing application of on-board processing is the extraction of information from high bandwidth video. An example of this has been accomplished for stationary video by the Motion Estimation and Image Chipping algorithms implemented in [4]. This process allows the node to transmit only what is moving/changing in the video (*imageChips*) rather than the whole video frame, greatly reducing bandwidth requirements. In addition, QED allows this information to be compressed, scaled, and/or decimated to match bandwidth resources.

A new algorithm, GMTI Segmentation, has been integrated that subscribes to GMTI tracks and *videoFrame* metadata (not the 4MB image payload) and computes which GMTI tracks fall within the bounds of the georectified image. Based on the geometry of the image and characteristics of the GMTI track, the algorithm generates a bounding box around each target within the frame and publishes this information as a *blobList*. Because GMTI Segmentation is publishing an existing type of information, *blobList*, it is interoperable with the Image Chipping algorithm and was easily integrated as a different means of extracting relevant information from high bandwidth video. The N-CET architecture makes information the interface between algorithms, permitting reconfiguration of the system by substituting algorithms to suit mission needs while being transparent to other algorithms in the system. For example, if a FLIR sensor was added to a node, one could envision an algorithm segmenting portions of a video frame above a threshold temperature and allowing the Image Chipping algorithm to extract that information.

C. *User Interfaces*

The clients described in the previous section reside on the node, next to the sensors. The information they produce is available to remote users through subscriptions, allowing a user to specify which types are relevant to one's mission and receive only those types. Several clients have been developed to allow users to generate these subscriptions, visualize the information, and interact with N-CET nodes.

The user interface referred to as the Console allows users to remotely connect to each node and establish subscriptions to desired MIO types. Predicates can also be specified to filter instances of MIOs within a type based on elements of the metadata, such as location or source.

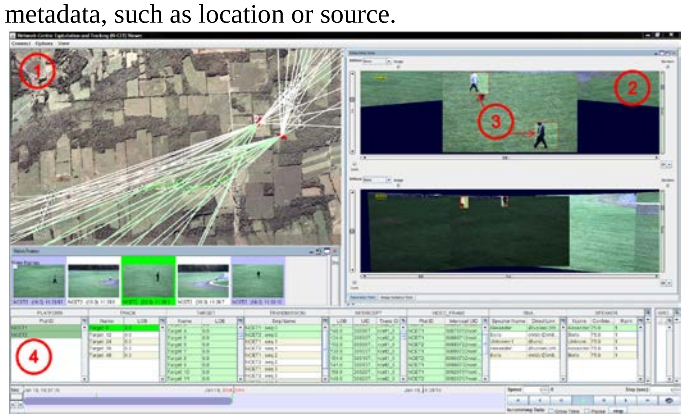

Fig. 1. Console User Interface

The Console (Fig. 1) has several panels that can be displayed and arranged by the user. The geospatial panel (1) is a three dimensional view that overlays geospatial information such as RF targets and GMTI tracks onto map and terrain data. A separate panel (2) displays information generated by the EO/FMV sensors. This includes the *imageChips* extracted from FMV (3) overlaid onto compressed background imagery. Other panels (not shown) include bandwidth status, and a control panel to publish user commands to selected nodes.

For testing and experimentation purposes, it is often necessary to subscribe to all the information being produced by multiple N-CET nodes. This large amount of information from disparate sources can be a challenge to visualize for the user, especially when the relationship between information must be portrayed. A major upgrade to the Console has been the development and integration of the Priority Filter Viewer (PFV). The PFV provides a means for a user to manipulate the visualization of information by ordering and selecting information to prioritize and filter, respectively, what is displayed. In the Console, information is displayed in a tabular format (Fig. 1 (4)) by type. A user may select a particular instance of information (a cell), such as the track of an RF emitter, and the other instances of information related to that track will be highlighted, such as the individual triangulated targets, imagery of the target, and any speaker identification or social network analysis results produced. The objects visualized in the geospatial panel are tied to the PFV and are highlighted as well. Double-clicking on an instance of information will hide all non-related information, and a type can be removed completely from view by removing its table. Information type tables may also be reordered to prioritize the type most relevant to the user. For example, moving the Speaker table to the leftmost position allows a user to focus on only the information related to a particular speaker.

The Console is more suited for use by an operator in the AOC supporting a JTAC rather than a warfighter in the field. However, light-weight clients have been developed for specific purposes, such as rendering imagery and image chips extracted from FMV, which could be hosted on tactical laptops and handheld devices. The availability of information in the N-CET architecture makes development of these tailored applications straight forward.

IV. Experimentation Results

The N-CET architecture has been evaluated through the integration of the sensors and algorithms described in Section III as well as field experimentation at the AFRL Stockbridge Test Site. The experiments demonstrated successful triangulation of RF emitters based on the LOBs generated by multiple platforms and cross-cueing of the EO/FMV sensor at each node onto these targets as well as GMTI tracks from remote (simulated) platforms. The architecture supported subscriptions from multiple users providing relevant information in a timely manner.

One aspect of the experiment evaluated the ability of the architecture to improve sensor resource allocation and utilization of bandwidth for relevant information. In an environment with multiple GMTI targets, an operator was interested in only cross-cueing the EO sensor onto targets in specific regions, called "watchboxes". The user was able to implement a subscription with a geospatial predicate so the Controller responsible for cross-cueing the camera would only receive GMTI targets within the watchbox. The results of the experiment are shown in Table 1. Without subscription-based filtering, roughly 4/5ths of the images were not relevant to user, meaning only 20% of the 49.8 MB of imagery sent to the user was pertinent to the mission objective. When filtering tracks based on a watchboxes, 100% of the 32 MB sent to the user was pertinent. The watchboxes limited sensor tasking to relevant targets and not only reduced the amount of non-relevant images sent onto the network but also increased the number of relevant images because sensor cycles were not wasted on unimportant targets.

TABLE 1
WATCHBOX EXPERIMENT RESULTS

	All Tracks	Filtered Tracks
Total Images	**159**	**104**
Images of Interest	**32**	**104**
Data Sent Offboard	**49.8 MB**	**32 MB**

The ability to extract moving objects from HD FMV was also demonstrated through experimentation. HD (1080p) video is collected at rate of 28 frames per second, producing 116 MB/s of raw video data. However, not all of this data is relevant and therefore does not immediately need to be sent offboard. Processing the video to extract moving objects greatly reduces offboard bandwidth requirements, but is dependent on the scene, i.e. the percentage of pixels in the frame that are changing.

Even at a reduced size compared to traditional FMV, *imageChips* can easily overwhelm the bandwidth capacity of a datalink, and potentially starve out other MIO types. QED, using the network visibility provided by VIA, manages the MIOs disseminated onto the network to provide QoS based on user defined policy. For example, a mission may dictate that *gmtiTracks* have a high priority. In this case, the dissemination of lower priority MIOs, such as *videoFrames* and *imageChips*, must be limited. Fig. 2 illustrates the performance of N-CET under such a policy. *gmtiTracks* are received at the Console frequently and with low latency, while *videoFrames* and *imageChips* have a much higher latency that grows over the duration of the scenario as these low priority MIOs queue up.

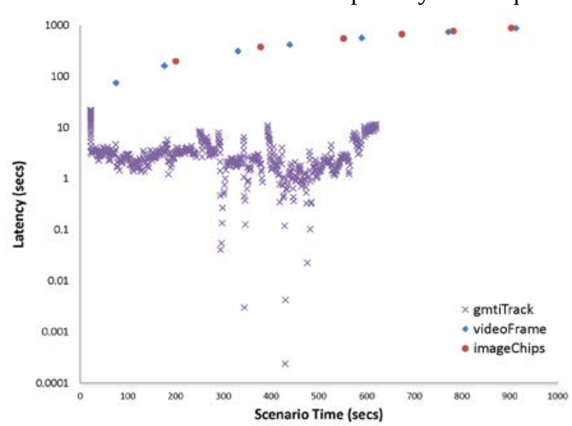

Fig. 2 MIO Arrival Latency. Arrival latency is measured from MIO creation time to arrival at Console.

When more information is produced than can be transmitted offboard, QED policy can be set to control what is disseminated. For the case above, the user receives all the information, but it can potentially be outdated, or stale. For near-real-time ISR applications, stale information has little value, so replacement policies may be employed in QED to ensure newer MIOs are favored over older MIOs. Fig. 3 shows the latency for two types of information, *videoFrames* and *imageChips*. In this scenario, a single compressed image is published followed by *imageChips* of moving objects extracted from FMV that will be displayed over the image, as was shown in Fig. 1 (3). The *imageChips* are generated at the same rate as the FMV, and the data rate far exceeds the observed capacity of the link (64kbps). Therefore, as new MIOs arrive, QED replaces older MIOs that are in the dissemination queue, reducing the latency of the MIOs and ensuring the user receives up-to-date information. Distinguishing between publication and arrival latency shows latency incurred as QED waits for the bandwidth capacity to send an MIO, and the latency of transmitting the MIO over the network, respectively.

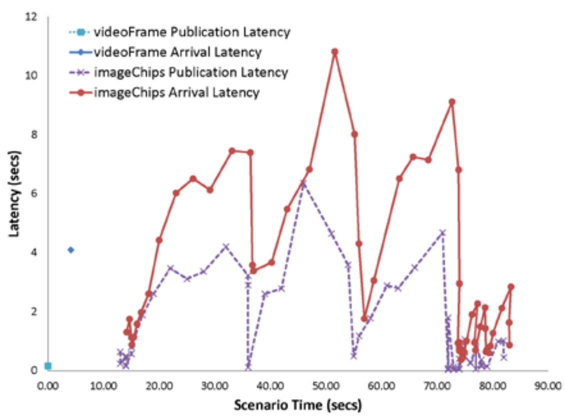

Fig. 3. MIO Publication and Arrival Latency. Publication latency is measured as the time between creation of the MIO and the dissemination of the MIO by the QED dissemination service.

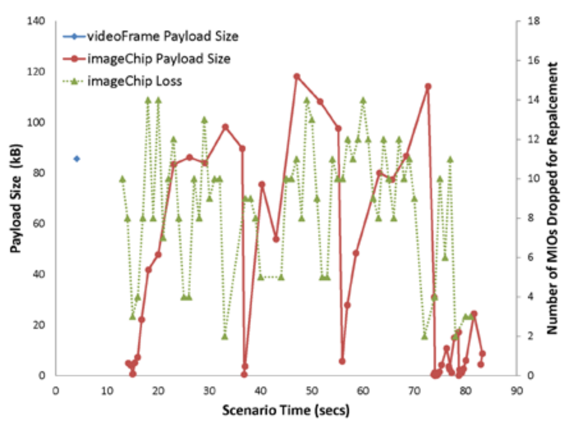

Fig. 4. MIO Payload Size and Loss. Loss is measured in the number of MIOs dropped by the QED dissemination service (right axis) in favor of replacement with a newer MIO.

The variation in the latency is the result of variations in the payload size of the MIOs. For *imageChips*, the payload size is dependent on the amount of motion in the scene. Payload sizes are shown in Fig. 4 for both the *videoFrame* and the *imageChips*. Also shown in Fig. 4 is the number of *imageChip* MIOs that are lost to replacement in order to disseminate the most recent data. No videoFrame MIOs are replaced. While these dropped MIOs are not disseminated to the user, the Phoenix IMS allow the MIOs to be archived on board, so that users can query for the full data at a later time if necessary, potentially on a higher bandwidth link.

In addition to motion-driven object extraction, GMTI tracks were successfully extracted from HD FMV. The accuracy of the blobs depended on the quality of the GMTI tracks and the accuracy of the position and orientation of the FMV sensor input into the georectification algorithm [10]. Improving those methods is beyond the scope of this work; however, the experiment demonstrated the correlation of information and processing at the sensor to allow real time extraction versus forensic analysis on the ground, as well the reconfigurability of the system to allow various video processing techniques.

V. Future Work

Research is already underway at the AFRL Information Directorate to integrate IM and embedded high performance computing into existing pods such as the LITENING AT. The architecture concepts developed and tested in N-CET will be transitioned to this program to provide tactical users contextualized, prioritized information from tactical assets.

Acknowledgment

The authors would like to thank the internal team members at the AFRL Information Directorate as well as the external participants for their contributions to this research project.

References

[1] E. Ghashghai, "Communications Networks to Support Integrated Intelligence, Surveillance, Reconnaisance, and Strike Operations," Rand Corporation, Santa Monica, 2004.

[2] C. P. Satterthwaite, "Testing Operational Fligth Programs (OFPs)," Defense Technical Information Center, 1994.

[3] Northrop Grumman, "AN/AAQ-28(V) LITENING Targeting Pod," [Online]. Available: http://www.es.northropgrumman.com/solutions/litening/.

[4] J. M. Metzler, M. H. Linderman and L. M. Seversky, "N-CET: Network-Centric Exploitation and Tracking," in *Military Communications Conference 2009. MILCOM 2009. IEEE*, Boston, 2009.

[5] R. Hillman, "Phoenix Architecture," Defense Technical Information Center, 2009.

[6] M. Slee, A. Agarwal and M. Kwiatkowski, "Thrift: Scalable Cross-Language Services Implementation," Facebook, [Online]. Available: http://thrift.apache.org/static/thrift-20070401.pdf.

[7] A. Mascarin and S. Thieret, "General-purpose GPUs breathe new life into high-performance embedded computing," VITA Technologies, [Online]. Available: http://vita-technologies.com/articles/general-purpose-life-high-performance-embedded-computing/.

[8] J. Loyall, *et al*., "Quality of Service in US Air Force Information Management Systems," in *MILCOM*, Boston, MA, 2009.

[9] M. Carvalho, *et al*., "Integrated Information and Network Management for End-to-End Quality of Service," in *MILCOM*, Baltimore, MD, 2011.

[10] T. Ager, *et al*., "Cameras and Sensing Systems," in *Manual of Photogrammetry*, Bethesda, American Society for Photogrammetry and Remote Sensing, 2004, pp. 581-677.